toulbar2
An exact cost function network solver

toulbar2 team.

Université de Toulouse, ANITI, INRAE, UR 875, 31326 Toulouse, France
contact author: simon.de-givry@inrae.fr

Mots-clés : *(probabilistic) graphical model, discrete optimization, constraint programming*

1 Introduction

toulbar2 is an open-source C++ solver for cost function networks, with an MIT license and a documentation describing its interfaces with C++ and python. [1]

The constraints and objective function are factorized in local functions on discrete variables. Each function returns a cost for any assignment of its variables. Constraints are represented as functions with costs in $\{0, \top\}$ where $\top$ is an upper bound cost associated with forbidden assignments. toulbar2 looks for a non-forbidden assignment of all variables that minimizes the sum of all functions. These functions can be nonlinear (quadratic terms or higher arity) and all variables must have a finite domain of reasonable size (less than a few hundred).

Using on the fly translation, toulbar2 can also directly solve optimization problems on other graphical models such as Maximum Probability Explanation on Bayesian networks, and Maximum A Posteriori on Markov random fields. It can also read partial weighted MaxSAT problems (*wcnf* format), pseudo-Boolean optimization problems (*opb* and *qpbo* formats), as well as constrained satisfaction and optimization problems (CSP and COP in XCSP3 format). [2]

toulbar2 exploits an *anytime* hybrid best-first branch-and-bound algorithm (HBFS) [1, 6] that tries to quickly provide good solutions together with an upper bound on the gap between the cost of each solution and the (unknown) optimal cost. Thus, even when it is unable to prove optimality, it will bound the quality of the solution provided. It can also apply a variable neighborhood search algorithm exploiting a problem decomposition (UDGVNS) [4]. Both algorithms are complete and can be run in parallel using OpenMPI [2, 4].

In addition, toulbar2 can also find a sequence of diverse solutions [5] or exhaustively enumerate solutions below a cost threshold and perform guaranteed approximate weighted counting of solutions. For stochastic graphical models, it computes the partition function, a #P-complete problem. Current work in toulbar2 is on multicriteria and bilevel optimization.

The presentation will give an overview of the solver and its results on two recent competitions.

2 XCSP'2022 Competition Results

The toulbar2 team participated to the XCSP'2022 competition. We used the default toulbar2 settings running HBFS for the *Main and Mini COP* tracks and parallel HBFS with 4 cores [2] for the *Parallel COP* track. toulbar2 terminates in second position in the Parallel and Mini COP tracks and forth position in the Main COP track (see Table 1 for Mini COP). [3] Among the five competitors in the Mini COP track (respectively 2 competitors in the Parallel COP track), it was able to solve the largest number of instances optimally (resp. 51/158 and 80/250).

1. https://github.com/toulbar2/toulbar2
2. https://xcsp.org
3. https://www.cril.univ-artois.fr/XCSP22

solver	score	optimum	best bound
Mistral	**93.00**	34	**99**
toulbar2	86.00	**51**	87
miniRBO	74.50	41	78
Sat4j-both	58.50	39	60
Sat4j-rs	43.00	33	46
Glasgow	31.50	21	34

TAB. 1 – XCSP'2022 Competition Results in the Mini Constraint Optimization Problem Track.

3 UAI'2022 Competition

We also participated to a competition on probabilistic graphical models. We compared our solver toulbar2 using two different methods, HBFS and UDGVNS, with cplex 20.1 and daoopt [3] on 120 instances provided as tuning benchmarks by the competition. [4] The average results are shown in Fig. 1. In 1 hour, toulbar2 solved optimally 86/120 instances whatever using HBFS or UDGVNS, whereas daoopt solved 92 instances and cplex 95. In less than two minutes, UDGVNS got the best solutions in average, while HBFS provided the best dual bounds.

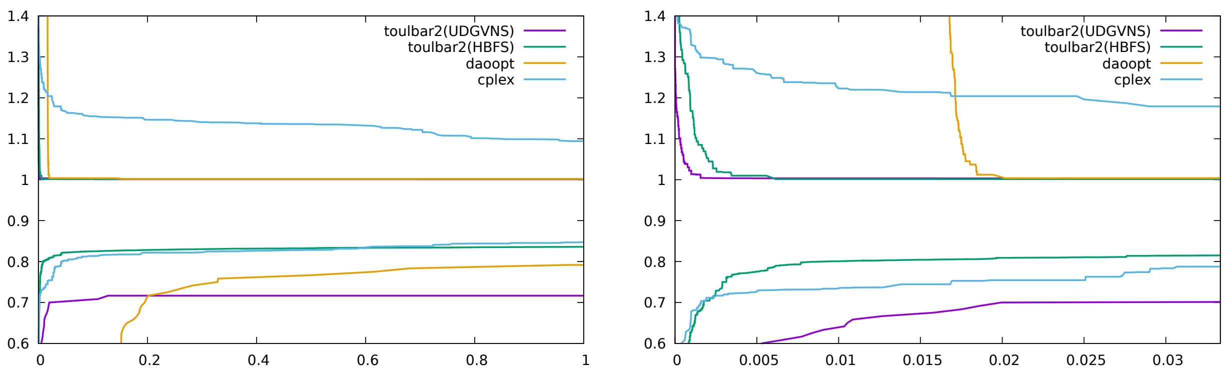

FIG. 1 – Average normalized lower and upper bounds (y-axis) as time passes (x-axis, Left : 1hour, right : 2min.) for cplex, daoopt, and toulbar2 on UAI'2022 tuning benchmarks.

Références

[1] D Allouche, S de Givry, G Katsirelos, T Schiex, and M Zytnicki. Anytime Hybrid Best-First Search with Tree Decomposition for Weighted CSP. In *Proc. of CP-15*, pages 12–28, Cork, Ireland, 2015.

[2] A Beldjilali, P Montalbano, D Allouche, G Katsirelos, and S de Givry. Parallel hybrid best-first search. In *Proc. of CP-22*, Haifa, Israel, 2022.

[3] L Otten, A Ihler, K Kask, and R Dechter. Winning the pascal 2011 map challenge with enhanced AND/OR branch-and-bound. In *DISCML'12 Workshop, at NIPS'12*, Lake Tahoe, NV, 2012.

[4] A Ouali, D Allouche, S de Givry, S Loudni, Y Lebbah, L Loukil, and P Boizumault. Variable neighborhood search for graphical model energy minimization. *Artif. Intelli.*, 278, 2020.

[5] M Ruffini, J Vucinic, S de Givry, G Katsirelos, S Barbe, and T Schiex. Guaranteed diversity and optimality in cost function network based computational protein design methods. *Algorithms*, 14(6) :168, 2021.

[6] F Trösser, S de Givry, and G Katsirelos. Relaxation-aware heuristics for exact optimization in graphical models. In *Proc. of CP-AI-OR'2020*, pages 475–491, Vienna, Austria, 2020.

4. `https://uaicompetition.github.io/uci-2022`